

49563



Decl Rule 132 OA2 S2_signed.doc

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re the Application of:	Application No.: 10/828,465
Fultheim, Shai	
Filed: 04/21/2004	Art Unit: 2128
For: Cluster Based Operating System-Agnostic Virtual Computing System	Examiner: Silver, David

DECLARATION UNDER 37 CFR 1.132

Mail Stop Amendment
Commissioner for Patents
PO Box 1450
Alexandria, VA 22313-1450

I, the undersigned, Dan Eylon, hereby declare as follows:

- 1) I am an employee of ScaleMP, the assignee hereof, in the capacity of Vice President for Research and Development (R&D).
- 2) I have worked in research and development and R&D management in the fields of computer software and algorithms for twenty years. I have received the following academic degrees: B.Sc., Physics and Mathematics, Hebrew University of Jerusalem (1984), M.Sc., Physics, Tel Aviv University (1988), and Ph.D., Physics, Tel Aviv University (1994).

I am familiar with the field of the documents Bugnion, *Disco: Running Commodity Operating Systems on Scalable Multiprocessors*: (Bugnion (Disco)) and Bugnion, U.S. Patent No. 6,075,938 (Bugnion (Patent)) (collectively referred to herein as "the Bugnion documents"), cited by the Examiner against the claims of the

Application hereof. Through my experience in management in the computer field, I am aware of the capabilities of those ordinarily skilled in the art of virtual computing and multiprocessors, and in particular in the field of the Bugnion documents.

3) The Bugnion documents disclose a single shared memory multiprocessor that is partitioned into several virtual machines.

The claims of the Application are directed to a cluster of machines having respective memories that are aggregated into a single large virtual computing system using multiple distinct virtual machine implementers, e.g., multiple virtual machine monitors.

The architectural differences between the systems disclosed in the Bugnion documents and the Application are apparent from a comparison of two Figures attached hereto as Appendix 1. Fig. 1 is a reproduction of Fig. 1 of Bugnion (Disco), cited by the Examiner, which shows the high level architecture of a single multiprocessor. Fig. 2 is an adaptation of Fig. 1 of the Application into a similar symbolic representation, and describes an embodiment of a cluster based virtual computing system disclosed in the Application.

There are major differences in scope and capabilities between the two systems. In the following discussion, concurrent references to the Drawings of the Application are given in parentheses for the convenience of the Examiner.

In the architecture of Applicant's claimed embodiments, shown in Fig. 2, the interconnect (Application Drawings, Fig. 1; network 44) is not part of the intrinsic hardware (Application Drawings, Fig. 1; PC Hardware 28) of the individual machines of the cluster, but is a separate element. The interconnect is used by each VM implementer (indicated in Fig. 2 by the symbol "VMM") (Application Drawings, Fig. 1; VMM 34, 36, 38) to form a single virtual machine

(Application Drawings Fig. 1; VM 20). The underlying processing elements (PE) (Application Drawings Fig. 1; nodes 22, 24, 26, PC hardware 28) are just separate, disconnected computers.

Disco as described in the Bugnion documents, assumes that the underlying hardware is cache coherent, so that the software layer 'sees' a shared memory hardware picture). As Bugnion (Disco) states (Sec. 4) (emphasis added):

"Disco is a virtual machine monitor designed for
the FLASH multiprocessor [17], a scalable cache-
coherent multiprocessor. The FLASH multiprocessor
consists of a collection of nodes each containing a
processor, main memory, and I/O devices. The nodes
are connected together with a high-performance
scalable interconnect. The machines use a directory
to maintain cache coherency, providing to the
software the view of a shared-memory multiprocessor
with non-uniform memory access times."

The architecture according to the claimed invention involves no such assumption. Rather, it implements cache coherency algorithms/mechanisms in software, using the interconnect to 'glue' several separate machines into one large virtual computer (Specification, paragraphs [0072], [0091] et seq.). The claimed embodiments, in sharp contrast with Disco, do not partition a large multiprocessor into several virtual machines, but does just the opposite: it takes several distinct standalone machines, and using multiple virtual machine implementers, aggregates the hardware of these machines into a single large virtual computer (Application Drawings, Fig. 1; PC Hardware 22, 24, 26, VMM 34, 36, 38, VM 20).

The Bugnion documents make no mention whatsoever of the use of multiple virtual machine implementers or monitors (VMM). On the

contrary, its VMM (indicated by the symbol "DISCO" in Fig. 1) always runs as a single entity directly above the hardware.

In the claimed invention, the single large virtual computer, aggregating separate, non-shared memory spaces into a single large shared memory. This, again, is in sharp contrast with the Disco implementation. Bugnion (Disco) states (Sec. 6.2):

"Disco is a virtual machine monitor that implements in software a virtual machine identical to the underlying hardware".

4) Based on the discussion above, it is hard to believe that the single paragraph in the conclusion of the Bugnion article (Sec. 7), referring to "more loosely coupled environments such as networks of workstations", refers to an implementation similar to virtual computer system as described and claimed herein. It makes much more sense to assume that the author meant that networking several machines, each running a Disco-like monitor, can enable simpler migration of processes from one workstation to another, simple checkpointing of process status, etc. The Bugnion documents do not mention aggregation of several separate machines into one virtual machine, implementing cache coherency over a network or an interconnect by software, turning several separate machine memories into one shared NUMA memory, or any other of the prerequisites required to build such a virtual machine.

5) A conversion of Disco into a virtual computer system as described and claimed in the Application would require major changes to the Disco architecture, as well as major additions to its content. (1) The virtual machine described and claimed in the Application does not "implement a virtual machine identical to the

underlying hardware", (2) The referenced distributed networks of workstations cannot "provide to the software the view of a shared-memory multiprocessor".

In fact, much of the complexity in the Disco design is relegated to the hardware, by assuming the underlying hardware is "ccNUMA", i.e., cache coherent. Cache coherency becomes extremely complicated to implement in software over a network-like interconnect. Moreover, the Disco system is described throughout the Bugnion documents as a single-instance VMM. This, of course, cannot hold true for separate computer systems combined into one virtual machine: as each computer system runs its own VMM, and uses the interconnect for its purposes, something that Disco does not do.

It is far from obvious, from the perspective of a person skilled in the art, what steps would need to be taken to transform the Disco system described in the Bugnion documents into a virtual computer system as described and claimed in the Application. The Bugnion documents fail to provide any guidance for modifying the disclosed Disco system in "more loosely coupled environments" to create a such a system.

6) That it is my considered opinion that:

(1) a person ordinarily skilled in the art would not interpret the Bugnion documents to refer to a cluster based virtual computing system as claimed herein;

(2) the disclosure of the Bugnion documents does not enable a person ordinarily skilled in the art to make the claimed invention; and

(3) it would not be obvious for a person ordinarily skilled in the art to modify the disclosed Disco system to make the claimed invention.

7) I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and conjecture are thought to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application of any patent issued thereon.

Dan Eylon

Dan Eylon,
Citizen of Israel

33 Harimon St.,
Ramat Hasharon 47251
Israel.

Date: March 11, 2008



Appendix 1

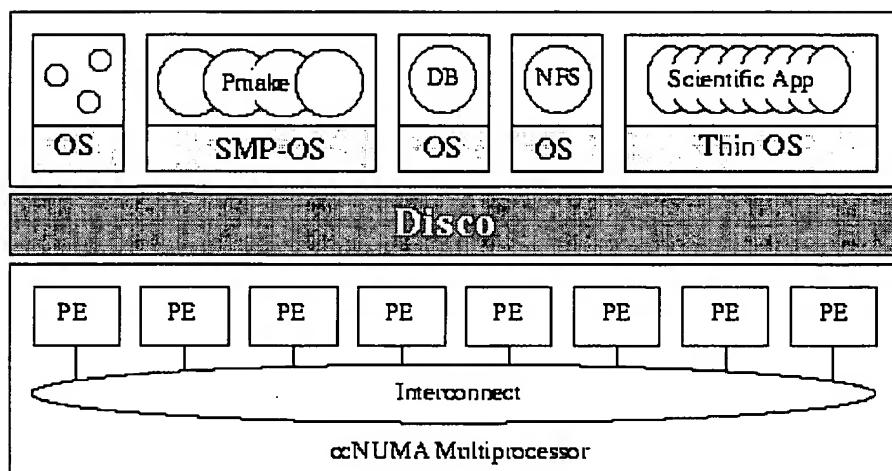


Figure 1

Source: Bugnion *et al.*

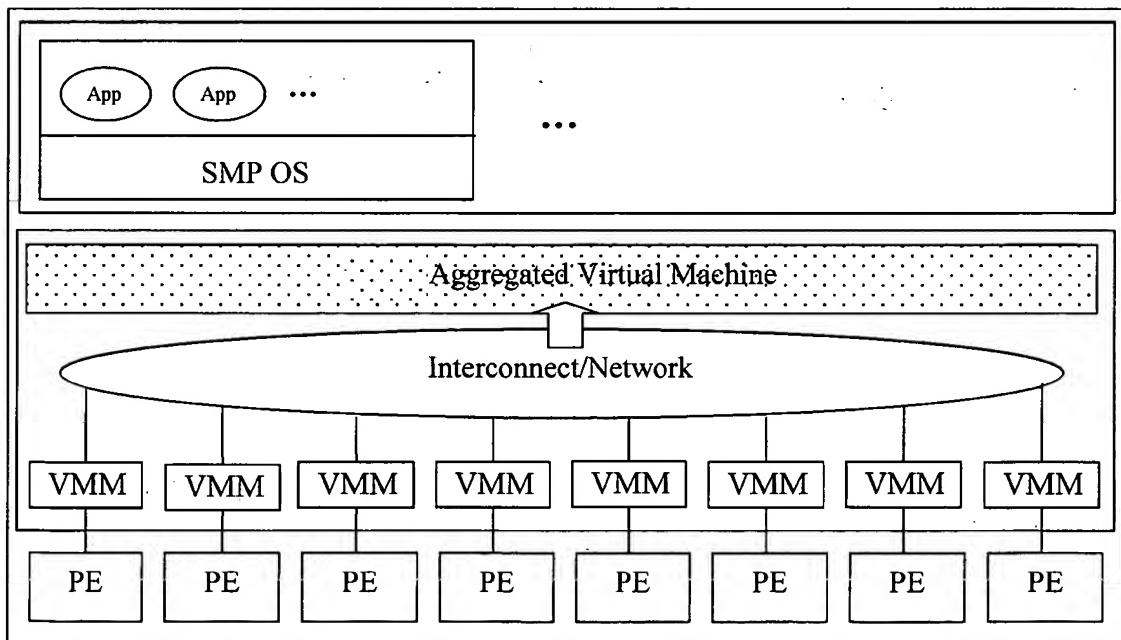


Figure 2